Doc ID: 6588666

(b) (3) - P.L. 86-36

NATIONAL SECURITY AGENCY
CENTRAL SECURITY SERVICE
FORT GEORGE G. MEADE, MARYLAND

P1 INFORMAL No. 3 JUNE 1978

An Application of Cluster Analysis to the Question of "Hands" and "Languages" in the Voynich Manuscript

S-216,867

By: M. E. D'IMPERIO
P13

Approved for Release by NSA on 08-03-2017, FOIA Case # 101982

Distribution:

[redacted]

[redacted]

(63) - P.L. 86-36

Doc ID: 6588666

I am reasonably certain that few readers of this paper will require much of an introduction to the topic of the Voynich Manuscript. Brigadier John Tiltman's informative and enjoyable presentation on 17 November 1975, and the seminar on 30 November 1976 served to familiarize many with this cryptanalytic challenge from the late Middle Ages. There have also been several articles on the subject in Cryptolog during the last few years. For any reader who desires an overview of the topic and a summary of some recent research, I recommend the Proceedings of our 1976 seminar [4] a copy of which may be obtained from M. DImperio, R53/P13. Two presentations by Captain Prescott H. Currier constituted high points of that occasion: in them, and in the supporting paper printed as Appendix A of the Proceedings, he set forth his theory that there were several different scribes involved in the production of the Voynich Manuscript, and that their individuality was attested not only by characteristic "hands", reliably distinguishable by eye, but also by statistically distinct "languages". If this hypothesis could be confirmed, it would provide students of the Voynich Manuscript with an important new insight into the problem. This paper describes the results of an exploratory study of Currier's theory using cluster analysis and multiple scaling. It has another purpose as well, primarily tutorial, in that I felt a detailed description of an application of these techniques to a relatively clear-cut problem might prove useful to others considering them for use in operational contexts.

Currier's Hypothesis

The Voynich Manuscript is a rather long document, comprising some 210 pages of writing in an unknown script liberally interspersed with colored drawings of a wide range of subjects and exhibiting (at least to us, today) a highly bizarre nature. The manuscript is considered to contain several sections, presumably dealing with different subject matter, as judged by the nature of the drawings. A long initial "herbal" section is profusely illustrated with representations of fanciful plants; an "astrological" section shows zodiacal diagrams and many illustrations featuring stars, suns, moons, and other cosmological elements; a "biological" section is marked by strange associations of naked female figures and objects like pipes, pools, and platforms; other sections are similarly distinguished by their illustrations. Currier's findings concern contrasts he has seen between sets of pages in certain sections of the manuscript, leading him to classify the pages into sub-groupings: an approach quite different from that of other students of the manuscript, who almost invariably consider it the monolithic production of one author.

Here are a few highlights drawn from Currier's exposition of his theory at the 1976 conference:

> "The first twenty-five folios in the herbal section are obviously in one hand and one 'language', which I call 'A'. ....The second twenty-five folios are in two hands, very obviously the work of at least two different men [(A and B]. In addition to this fact, the text of this second portion of the herbal section (that is, the next twenty-five or thirty folios) is in two 'languages' [A and B], and each 'language' is in its own hand. This means that, there being two authors of the second part of the herbal section, each one wrote in his own 'language' ....Now with this information available, I went through the rest of the manuscript...and in four other places I discovered the same phenomena I associated with 'language' B....The biological section is all in one 'language' (B) and one hand." [4, p.20ff.]

While he finds indications of different hands and "languages" in other sections of the manuscript (the pharmaceutical, astrological, and "recipe" sections), these seem much less distinct and clear-cut. It should be noted that in using the word "language" in this context, Currier does not necessarily mean to imply that he has found different underlying natural languages (e.g., Greek as against Latin, or German as against French). He is referring to patterns of statistical characteristics that seem to be consistently associated with Hand A as opposed to Hand B: certain symbols are more likely to occur together or to appear more frequently in certain positions in the "words" of the Voynich text in folios showing one hand than in folios showing the other. An inspection of his extensive monographic, digraphic, and trigraphic counts, and his studies of symbol clusters in various positions of a "word", has convinced him of the presence of at least two clearly distinct bodies of text. In these two corpora the symbols show certain consistently different and characteristic distributions, associated with the visible differences in writing style and formation of symbols marking the hands of two different writers (writer A and writer B). Currier refers to the two bodies of text as "languages" A and B. In summing up his findings, he indicates that he feels quite certain of at least five, and perhaps as many as eight, different hands in the manuscript as a whole, but only two statistical "languages".

This, then, is the exciting hypothesis put forward by Currier. Several of us, after attending his presentation, confirmed his suggestions to our own satisfaction by replicating his original procedure of choosing some pages showing obviously different writing styles in the large herbal section (where the contrasts between scribes A and B are especially striking), and verifying both the consistent differences in hand and certain clear accompanying differences in symbol patterns. Nevertheless, since so many other approaches to the problem posed by the Voynich Manuscript have been fraught with subjectivity and self-delusion, it seemed important to place Currier's findings on a more objectively demonstrable and secure basis, and to attempt to confirm or disconfirm them by an independent statistical study.

Cluster Analysis

In considering statistical tools for investigating Currier's hypothesis, I decided upon that of cluster analysis as an appropriate method. Cluster analysis algorithms are available as computer programs and are widely employed in the social and natural sciences for classifying collections of objects into subsets based on similarities and dissimilarities with respect to a list of scores or observations. The methods can also be used to reveal which of a group of objects is most like another single object in the group. So long as a set of observations has been made, such that every object under study has been scored, rated, or labelled for all the same properties or "variables", the clustering techniques can be applied to reveal subgroups among the objects. Within each cluster, objects are more like each other than like objects in other clusters. This methodology seemed to me a good choice for revealing the sets of similar pages within the Voynich manuscript that Currier's theory called for, if in fact they were present. For the reader interested in knowing more about it, a number of more or less readable works are available. Cluster analysis has been investigated by RSI for possible applications to Agency problems, and two excellent survey papers by Douglas A. Cope provide a summary of various clustering algorithms [3], and multiple scaling and related techniques [2]. A number of good reference works are available in the open literature; two I found particularly useful were Everitt 1974 [5] and Anderberg 1973 [1].

Doc ID: 6588666

There are numerous ways of carrying out cluster analysis, and the published computer programs embody various combinations of these, considered by their designers to offer some special advantage for certain applications. In general, however, the analysis involves the following stages: 1) deciding upon a group of objects that constitutes a good sample of the groupings or clusters hypothesized by the analyst; 2) deciding upon the observations to be made across all the objects; 3) taking the measurements, scores, rankings, labellings, etc. of each object with respect to each observation; 4) choosing a measure of "distance" (dissimilarity) or, alternatively, a measure of association (similarity) appropriate to the case; 5) computing the distances (or associations) between each object and every other with respect to the observations; and, finally, 6) applying the clustering algorithm to the triangular matrix of distances or associations resulting from step 5. The clustering procedures may be "agglomerative" (beginning with one object and iteratively joining other objects to it to form a cluster, as if crystalizing around a nucleus), or "divisive" (starting with all the objects in one big group and successively splitting them into dissimilar subgroups until no further splits can be made). Within the agglomerative methods, there are further choices among methods of linking new objects to existing clusters: "single linkage" methods focus upon the dissimilarity between nearest neighbors in a cluster, "complete linkage" methods upon the dissimilarity between the farthest neighbors, and "average linkage" methods upon the average dissimilarity among neighbors. As might be imagined, some methods are better at stringing out objects in long, thin chains, while others are better at dealing with globular clumps.

Unavoidably, as is the case with many sophisticated statistical tools, there is a real danger of imposing spurious structure upon the data if the techniques are badly chosen or unintelligently used. A factor analysis or multiple correlation method will find "factors" of some sort in any data, however difficult they may be to interpret or put to use. Similarly, a cluster analysis will always find clusters, and it is up to us to pay attention to the indications of significance (the relative compactness of the clusters, the strength of their internal "bonds", and the relative distance between clusters) as shown by the statistical measures which the programs provide as a part of the printout. The interpretation of cluster analysis results is unavoidably circular; we propose a certain structure in the group of objects under study, we perform the computation, and we are happy if we see what we expected, or at least something that makes sense in terms of our original hypothesis, however revised. If the clusters we get bear little or no relation to any groupings we expected to see, and we can make little sense of them in the context of our understanding of the problem, we have some indication that our hypothesis about clusters in the data was not confirmed, but an attempt to reason from such unexpected and apparently meaningless structures backwards to the data may prove unrewarding.

I will not attempt here to go into the details of the cluster analysis algorithms or the various methods of computation; I urge the interested reader to consult the references mentioned above. Instead, I will provide some varied examples of applications in which cluster analysis has proven useful, as a means of communicating the "flavor" of these methods. A frequent use of cluster analysis is in studying the genetic similarities among species of plants or animals, based on some set of chemical or physiological properties. Cluster analysis has been employed in statistical pattern recognition, to discriminate printed letters, geometric shapes, or other visual forms. In archaeology, it has been used to classify groups of artifacts gathered by surface collection over a site; clusters of similar objects concentrated in certain areas within the site were found to indicate different human activities ("women's" activities: cooking, spinning, making pots, vs "mens'" activities: weapons manufacture, hunting, herding). A particularly interesting application of cluster analysis to Egyptian Archaeology, published in a recent issue of Science, deserves special mention [6]. I will describe it at somewhat more length, since it demonstrates so dramatically the usefulness and power of this methodology when intelligently employed.

In 1898, a large cache of Egyptian royal mummies was found in the Valley of the Kings; these mummies, having been plundered and damaged by tomb robbers, had then been gathered together by a later Egyptian ruler, rewrapped, and deposited in two new hiding places. In the process of reburial, the identities of certain mummies were obscured (at least for the modern archaeologist). One in particular, referred to by archaeologists as the "elder lady", or more objectively, as Egyptian Museum Catalog Number 61070, was particularly interesting since it appeared from certain evidence (e.g., the position of the hands) to be that of a queen. It was suggested that this lady might be Queen Hatshepsut or Queen Tiye (mother of the heretical pharaoh Akhenaton). A set of coordinated studies were undertaken, including data from conventional full-body x-rays, standardized x-rays of the head known as "cephalograms", and scanning electron microprobe comparison of hair samples known to be from Queen Tiye and those from the unidentified mummy. Several different cluster analysis algorithms were applied to sets of scores obtained from cephalogram studies of the unknown lady and ten other mummies of Egyptian queens. The analysis showed clearly that the head measurements of the unknown matched those of Queen Tiye's mother more closely than those of any other queen. This finding was strongly supported by the close match between the hair samples known to belong to Queen Tiye (and obtained from a keepsake in the tomb of another family member) and hair from the unidentified "elder lady".

Application to Voynich Manuscript Pages

Selecting the objects. I was fortunate enough to have at my disposal a large corpus of text from the "herbal" and "biological" sections of the manuscript, transcribed according to the alphabet designed by Currier for computer processing of the Voynich symbols. Currier stated that he had found no page to be broken by a change of hand or "language", so that a set of samples, each taken from the text of a single page, should provide an appropriate test of Currier's theory. I selected forty segments of text, consisting of the first 350 to 400 characters from each of forty different pages. According to Currier's view, the text of these pages should fall into three major classes: herbal pages in "language" A and hand A, herbal pages in "language" B and hand B, and biological pages in "language" B and hand Z. These three classes will be called Herbal A, Herbal B, and Biological B for short in the remainder of this paper. Figure 6 shows a summary of pages from which samples were chosen.

Making the Observations. I decided upon a simple monographic frequency count as a good starting point, since Currier had found a clear difference in the distribution of individual symbols between "languages" A and B. I made forty monographic distributions, one for each of the selected pages, including roughly the first 350 to 400 characters on each sample page (many pages did not contain more than 400 characters, and I wished the samples to be more or less equal in size).

Choosing a Measure of Association. Since my data consisted of frequency counts applied to a set of mutually exclusive, exhaustive events (the symbols of the Voynich script "alphabet") I could avoid the many scaling and normalization problems afflicting investigators employing cluster analysis for sets of observations comprising disparate measurements. My frequency counts constituted a set of discrete, countably infinite or finite variables, on a scale having a zero point and permitting proportional measurement (i.e., if xi and xj are two counts within one distribution, we can say that xj is n times as large as xi). Therefore, I could consider my analysis to involve a "ratio" scale, the strongest of the four possible scales (ratio, interval, ordinal, and nominal) on which observations can be made. This left me free to use a wide variety of cluster analysis programs, employing various association measures. The analysis takes place in the context of a sort of abstract "measurement space" or "metric", within which the objects (manuscript pages represented by frequency distributions) are "located" at various "distances" from each other to form the clusters. Different programs may use any of several possible association measures, among them the Euclidean distance measure, the "city-block" distance measure (both measures of distance, or dissimilarity), and the correlation coefficient (a measure of similarity).

Computing the Association Matrix and Clusters. Through the courtesy of Douglas A. Cope, R51, I was able to obtain runs of four different cluster analysis programs and one program for multiple scaling. These programs were as follows: the PEP-1 Guttman-Lingoes Graph Theoretic Clustering algorithm; HICLUS (Hierarchical Clustering), an agglomerative method using single and complete linkage; TAXMAP-2, an average and single-linkage approach to mode-seeking; Ling's (K,R)-clustering Algorithm, a hierarchical k-linkage method; and MINISSA (Minnesota-Israel-Netherlands Integrated Smallest Space Analysis), the multiple scaling program. These programs have all been adapted for the CDC-6600 computer by Mr. Cope and his colleagues, and are described in his papers [2, 3]. His ingenuity and helpfulness to users in applying the techniques to their problems and interpreting the results have also been a major asset.

I needed only to supply a hypothesis and the set of forty frequency distributions, and the programs then carried out all the computations of associations, finding the clusters, and providing statistical estimates of confidence for the strength of the clusters or the program's representation of the data. In general, each of the cluster analysis programs found a lower triangular matrix of associations (correlation coefficients in PEP-1, HICLUS, and the K,R Algorithm, and city-block distance in TAXMAP). Each association in the matrix measured the relation of one Voynich Manuscript page, as represented by its monographic counts, to another single page. Transformations were then applied iteratively to the rows and columns of this matrix so as to emphasize the similarities and differences between pages. In some cases, the programs actually shuffled the rows and columns to bring like objects closer together in a final output matrix display; this was true of the (K,R) algorithm. As each cluster was found, a confidence measure was computed and associated with it in the program output as an aid to interpretation. The MINISSA program employed a somewhat different statistical model of the data; instead of finding clusters of objects in an abstract "space", it mapped the "locations" of the objects within such a space: a "Euclidean metric space", whose two dimensions may be assigned a meaning in relation to the hypothesis held by the investigator.

Interpreting the Results. Many programs provide a helpful graph or plot of the numerical results; in some cases, additional programs can be run on the outputs of a clustering algorithm to rearrange matrix rows and columns or provide graphic displays to aid the researcher. These visual representations are extremely helpful, and I found them almost necessary; unadorned lists of cluster members, ranged in dense rows down the pages of printout, can prove tedious and confusing indeed to the researcher. Since both the clustering and multiple scaling techniques are essentially applying a spatial model or "metaphor" to the problem posed by the investigator, a two-dimensional graph or plot is often an appropriate display. Another useful display is a "tree" or "dendrogram" showing the familial relationships among the objects. Each program provides statistical measures, associated with the clusters, the nodes of a dendrogram where each cluster is split off, or with the entire representation of the data. These measures are intended to enable the researcher to assess the confidence he may have in the findings of the program. In the next section, the outputs of the five programs will be described in detail.

Results of the Analysis

1. PEP-1 Graph Theoretic Algorithm. PEP-1 provides a list of clusters in order as each subset of the objects is partitioned off from the rest. A "family tree" (shown in Figure 1) can be drawn from this output. At each node of the tree where a cluster or a single object branches off, an "edge connectivity probability" is shown; this is an estimate of the likelihood that the split could have happened by chance. Thus, the lower this estimate of probability (on a scale of 0.0 to 1.0), the more confidence we may have in the contrast of the pages in the cluster against the rest of the pages outside of it. The upper "stem" of the family tree shows a loose sequence of small clusters and isolated pages, all from the Herbal A pages except for two samples, TL and HD (pages 94 and 76), from Herbal B. The tree then separates into two main branches; the left branch seems to correspond roughly to Currier's "language" B, since it contains most of the B pages and none of the A pages; all the Biological B pages are clustered together at the lower left, along with one Herbal B page (59); another Herbal B page (79) is alone, and there is another cluster of seven Herbal B pages just above. The right branch contains the rest of the Herbal A pages and one oddball Herbal B page (sample TE, page 60). The "probability" statistics seem quite low everywhere except in the right branch, where they suddenly jump up from near zero to .5. Thus, this right branch, while strongly split away from the rest of the tree, seems very weakly subdivided, and should probably be regarded as one very diffuse cluster.

2. Ling's (K,R) Algorithm. Figure 2 shows two triangular matrices output by the program. The rows and columns of an original "similarity matrix" containing correlation coefficients have been rearranged to place similar manuscript pages closer together and dissimilar pages further apart. Symbols made up of overstrikes were printed in the cells of the matrix, so that the higher correlations are darker; thus the clusters showed up as darker triangles along the main diagonal of the big matrix. A matrix was produced for each of several "bond sizes" or values of a threshold K applied to links between objects in clusters. Higher levels of K represent more restriction

on clusters, and a requirement for more strongly bonded clusters. Thus, for bond size 1 (k = 1), every object in a cluster must be joined to at least one other object in the same cluster by a link of the required closeness, and clusters must have at least two members. For bond sizes 2 and 3, every object must be linked to at least two or three others in the same cluster, and clusters must have at least three, or four, members. While higher bond sizes could have been required, the algorithm produced three matrices, one for each of bond sizes 1, 2, and 3. The first two matrices were essentially alike, and were as shown in the left drawing of Figure 2. Except for four anomalous Herbal B pages (59, 60, 76, and 94), there appear to be three relatively clear clusters corresponding to the three classes of pages Currier's theory calls for: Biological B, Herbal B, and Herbal A. The matrix for bond size 3 is somewhat different; it seems clearly to show only two major clusters, corresponding to Currier's two "languages", with the exception of the three Herbal B pages (60, 76, 94) and one Biological B page mixed in with Herbal A.

HICLUS Agglomerative Cluster Analysis. The output of HICLUS includes a dendrogram in the form of a display similar to a bar graph. Boundaries between clusters can be seen where low points in the graph-like display leave deep columns of white paper between the relative peaks of the clusters. An accompanying page associates sigmages for cluster tightness to each cluster; the higher the statistic, the tighter the cluster. The vertical dimension of the graph shows descending correlation coefficient values, so that objects associated in a cluster at the top have higher correlations, while the correlations decrease down the page. Figure 4 shows a rough redrawing of this bar graph. We see a rather strong cluster on the far right containing seven Herbal B pages; in fact, they are the same seven as appeared in the 7-page Herbal B cluster on the left branch in the PEP family tree. It has a correlation coefficient level of no less than .959, and a sigmage of 5.3. In the middle is the Biological B cluster, containing all the Biological B pages plus two from Herbal B (59 and 79); its correlation level without page 79 is .960, and its sigmage 7.97; with page 79, the corresponding figures are .947 and 8.47. The left half of the graph tails off into a very loose conglomeration of small clusters comprising all the Herbal A pages with two Herbal B pages (60 and 76). Page 94 is alone as an "outlier" (an object not clustered with any other in the set) on the far right.

TAXMAP-2 Clustering Program. While TAXMAP does not provide a graphic display of its results, Mr. Cope kindly ran its outputs through another program to create a two-dimensional "vector plot" similar in appearance to that produced by the MINISSA program discussed below. It should be noted at the outset that TAXMAP, alone of the programs run on my data, did not employ a correlation coefficient as a measure of association. Instead, a very different kind of measure was used: the "city-block" distance. In effect, this means that much of the information in my frequency count data on a ratio scale was disregarded; instead of comparing the profiles of peaks and valleys along the frequency distributions, a much cruder, less sensitive, and perhaps less appropriate measure of distance was used. This consideration may help to explain the differences in the results of TAXMAP as contrasted with those of all the other programs. The only cluster that shows up at all clearly contains ten Biological B pages. The Herbal A and B pages are scattered among small clusters and isolated individual pages in a manner that tells us relatively little that is useful.

MINISSA Multidimensional Scaling Program. Figure 3 shows a drawing adapted from the "vector plot" produced by the MINISSA program. Whatever feature of the metric space is represented on the horizontal axis, it seems to be related to the differences between Currier's "languages" A and B. The vertical axis is somewhat more problematical; it could pertain to the "subject" difference between Herbal abd Biological, or even to the "hand" difference between Currier's hands B and Z within "language" B. All the Biological B pages fall within a small, compact region in the lower center which also contains one anomalous Herbal B page, 59. A compact region above contains the same seven Herbal B pages we have seen grouped in a strong cluster by three of the four clustering algorithms. Page 94 is all alone on the extreme right, and page 76 is alone at some distance above and to the left of center, while pages 60 and 79 are around the edges of the Biological B region. All the Herbal A pages are scattered loosely over the leftmost third of the plot. In general, this program, in spite of its reliance on a somewhat different statistical model, appears to confirm the findings of PEP, the (K,R) Algorithm, and HICLUS.

Conclusions

Figure 5 shows a rough summary of the groupings of pages found by the five programs. Except for TAXMAP, they all seem to reveal the same picture: A strong Biological B cluster including all the Biological B pages along with one Herbal B page, 59; another cluster containing seven or eight of the Herbal B pages; and a loose association of Herbal A pages mixed with the same few anomalous fugitives from Herbal B. The transcribed data at my disposal do not contain a broad enough sampling from all sections of the manuscript to support a full-scale analysis attempting to study all the "hand", "subject" and "language" contrasts. Nevertheless, the results of this exploratory study clearly seem to be sufficiently encouraging to warrant a more complete analysis when more text has been transcribed.

Bibliography

1. Anderberg, M.R., Cluster Analysis for Applications," Academic Press, New York, 1973

2. Cope, D.A., "A Brief Introduction to Multidimensional Scaling and Smallest Space Analysis," R51, 1978

3. Cope, D.A., "Cluster Analysis: Introduction to Models and Methods," S-215,924, R51/MATH/25/76, 16 November 1976

4. DImperio, M.E., "New Research on the Voynich Manuscript: Proceedings of a Seminar," Washington DC, 30 November 1976

5. Everitt, B.S., Cluster Analysis, Halsted Press, 1974

6 Harris, J.E., et al., "Mummy of the 'Elder Lady' in the Tomb of Amenhotep II: Egyptian Museum Catalog Number 61070," Science, vol. 200, 9 June 1978, pp. 1149-1151

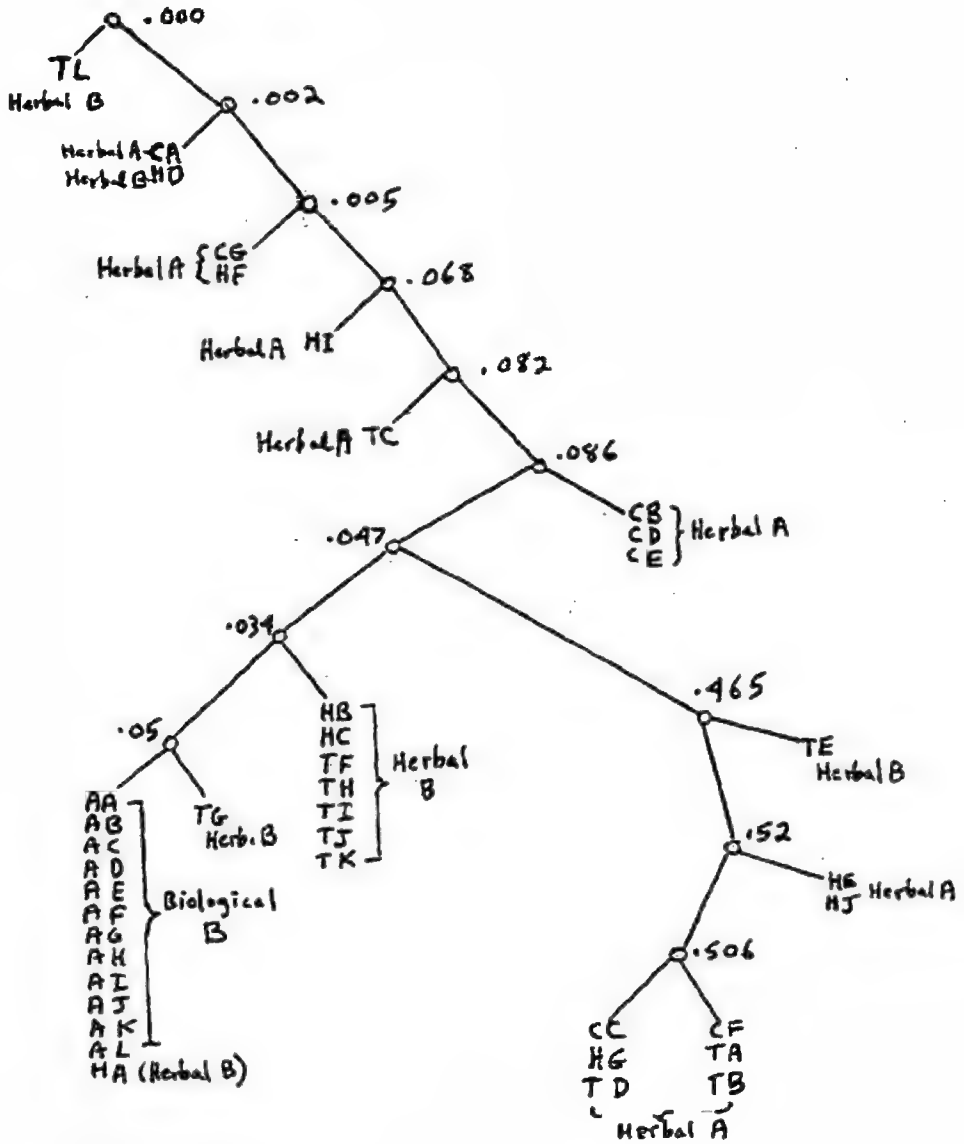

PEP-1 "Family Tree" (Single Linkage Clusters)

Figure 1

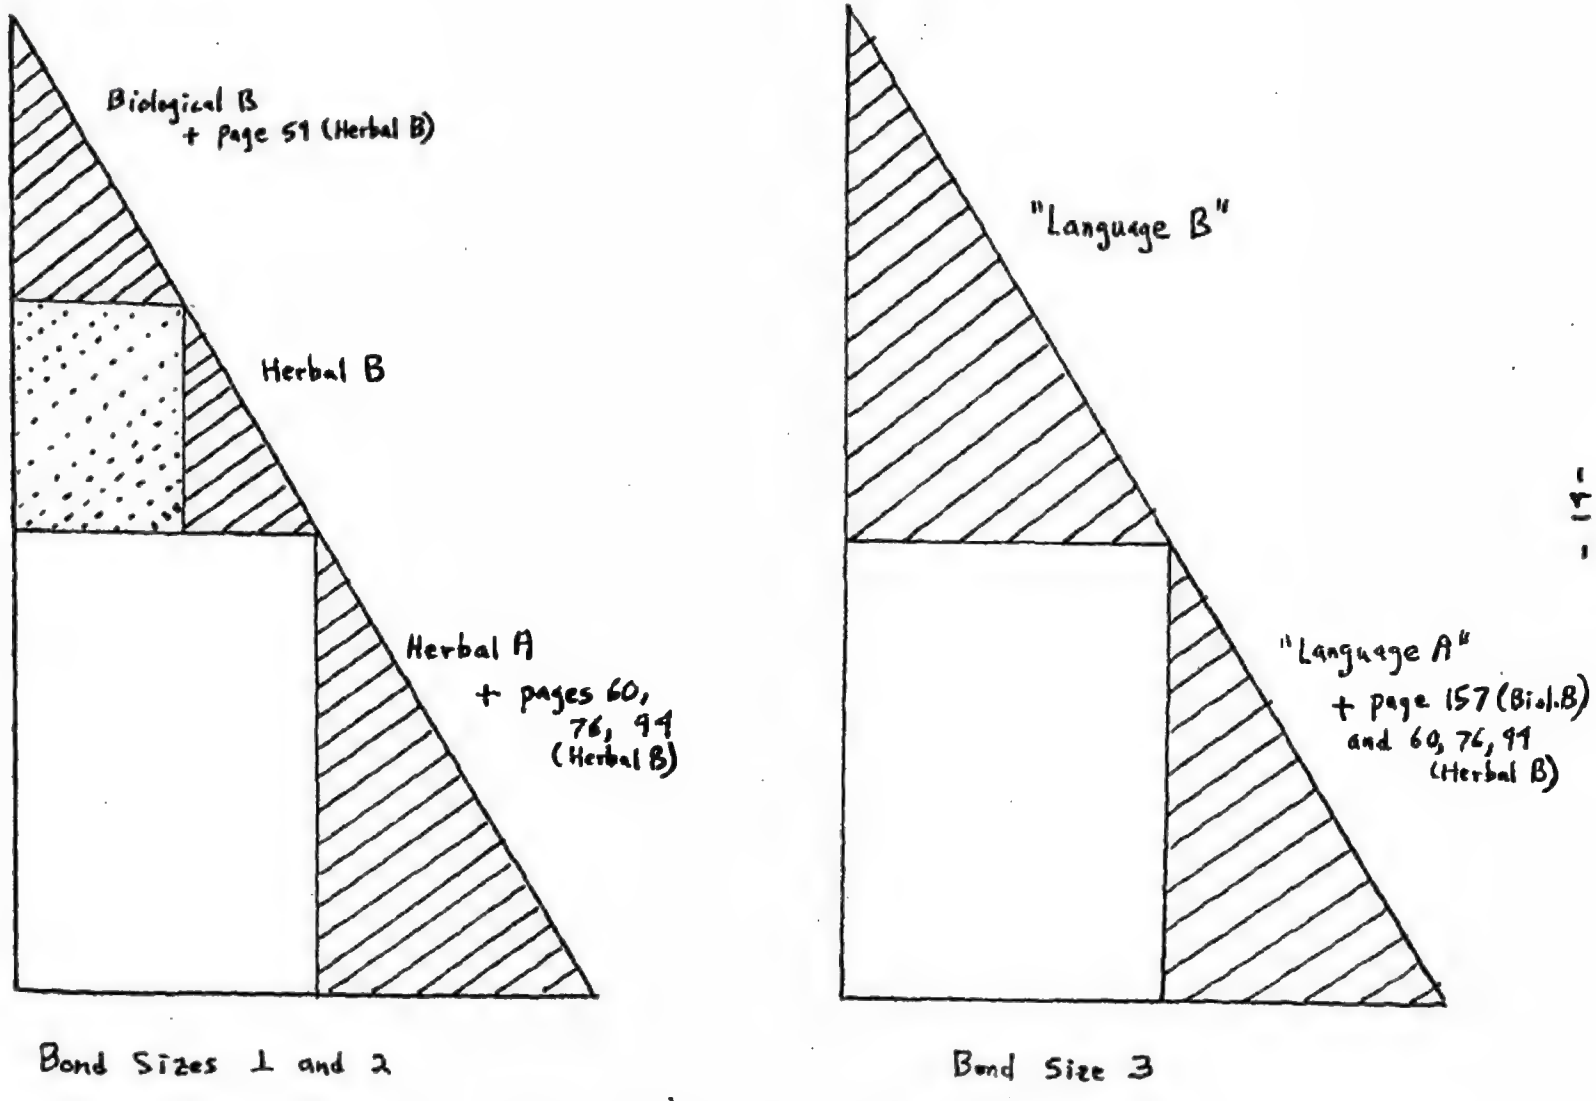
Biological B
+ page 51 (Herbal B)
Herbal B
Herbal A
+ pages 60, 76, 94 (Herbal B)
"Language B"
"Language A"
+ page 157 (Biol.B) and 60, 76, 94 (Herbal B)
Bond Sizes 1 and 2
Bond Size 3

Results of Ling's (K,R) Clustering Algorithm

Figure 2

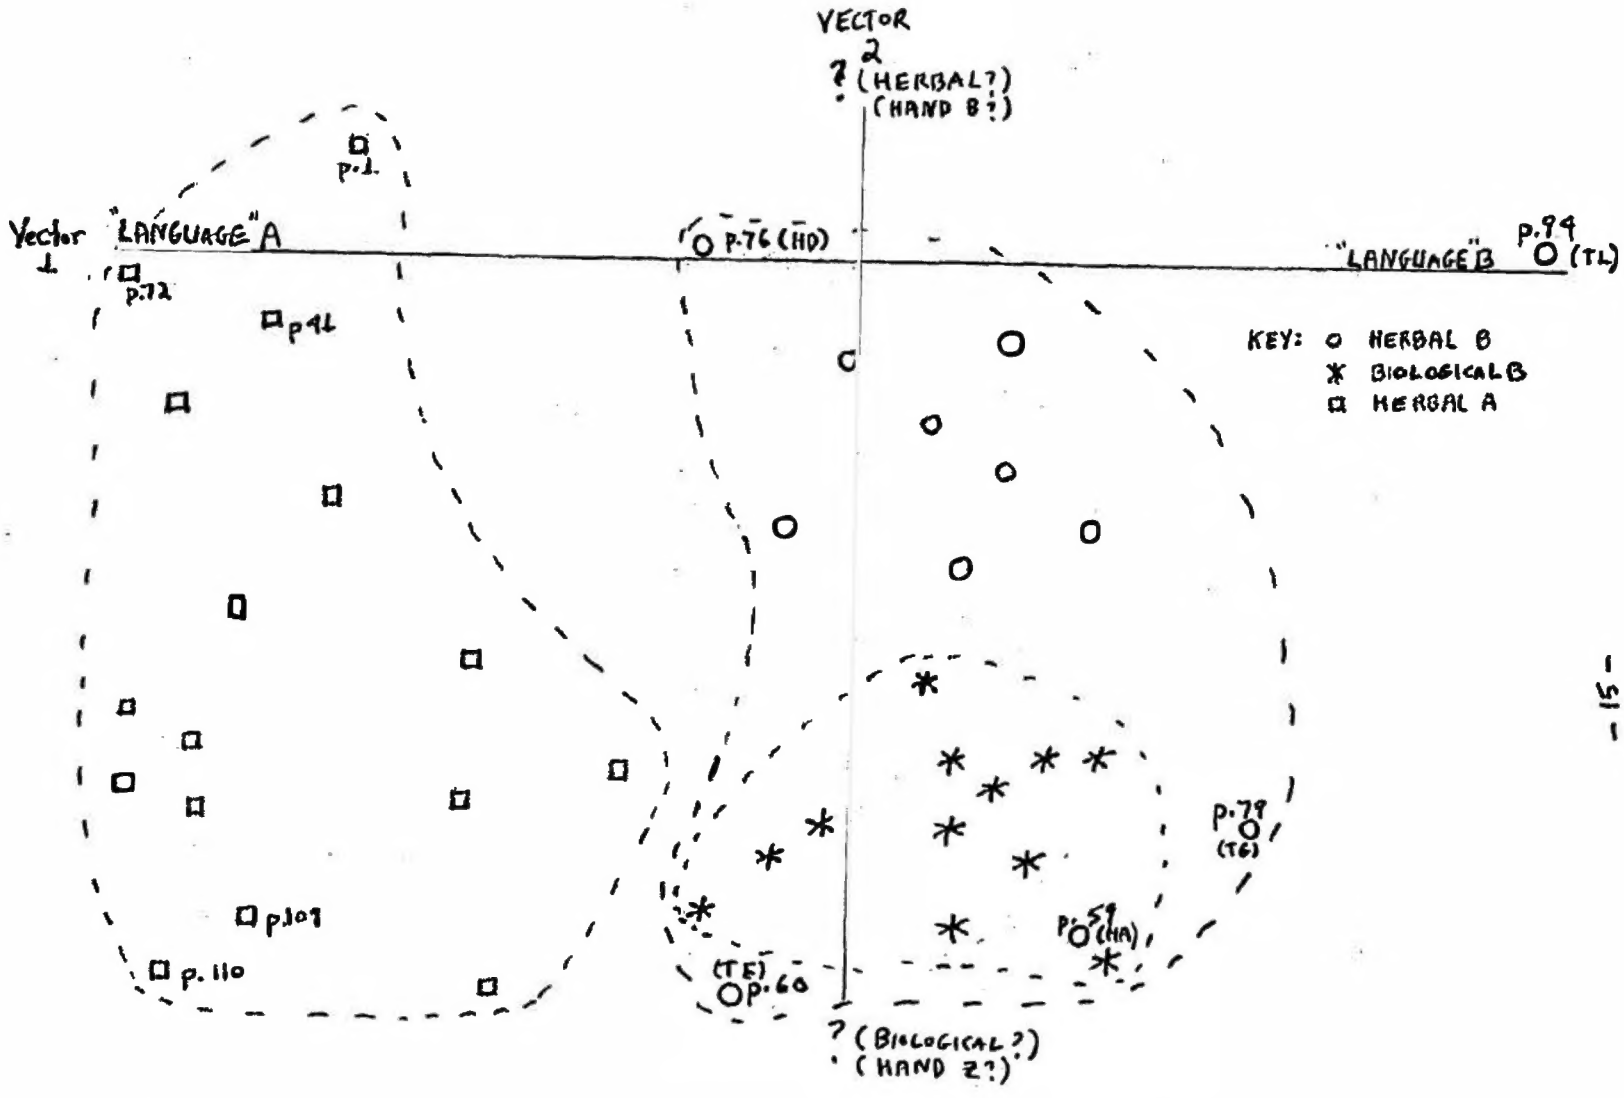

Results of Minissa (Multidimensional scaling) Program

Figure 3

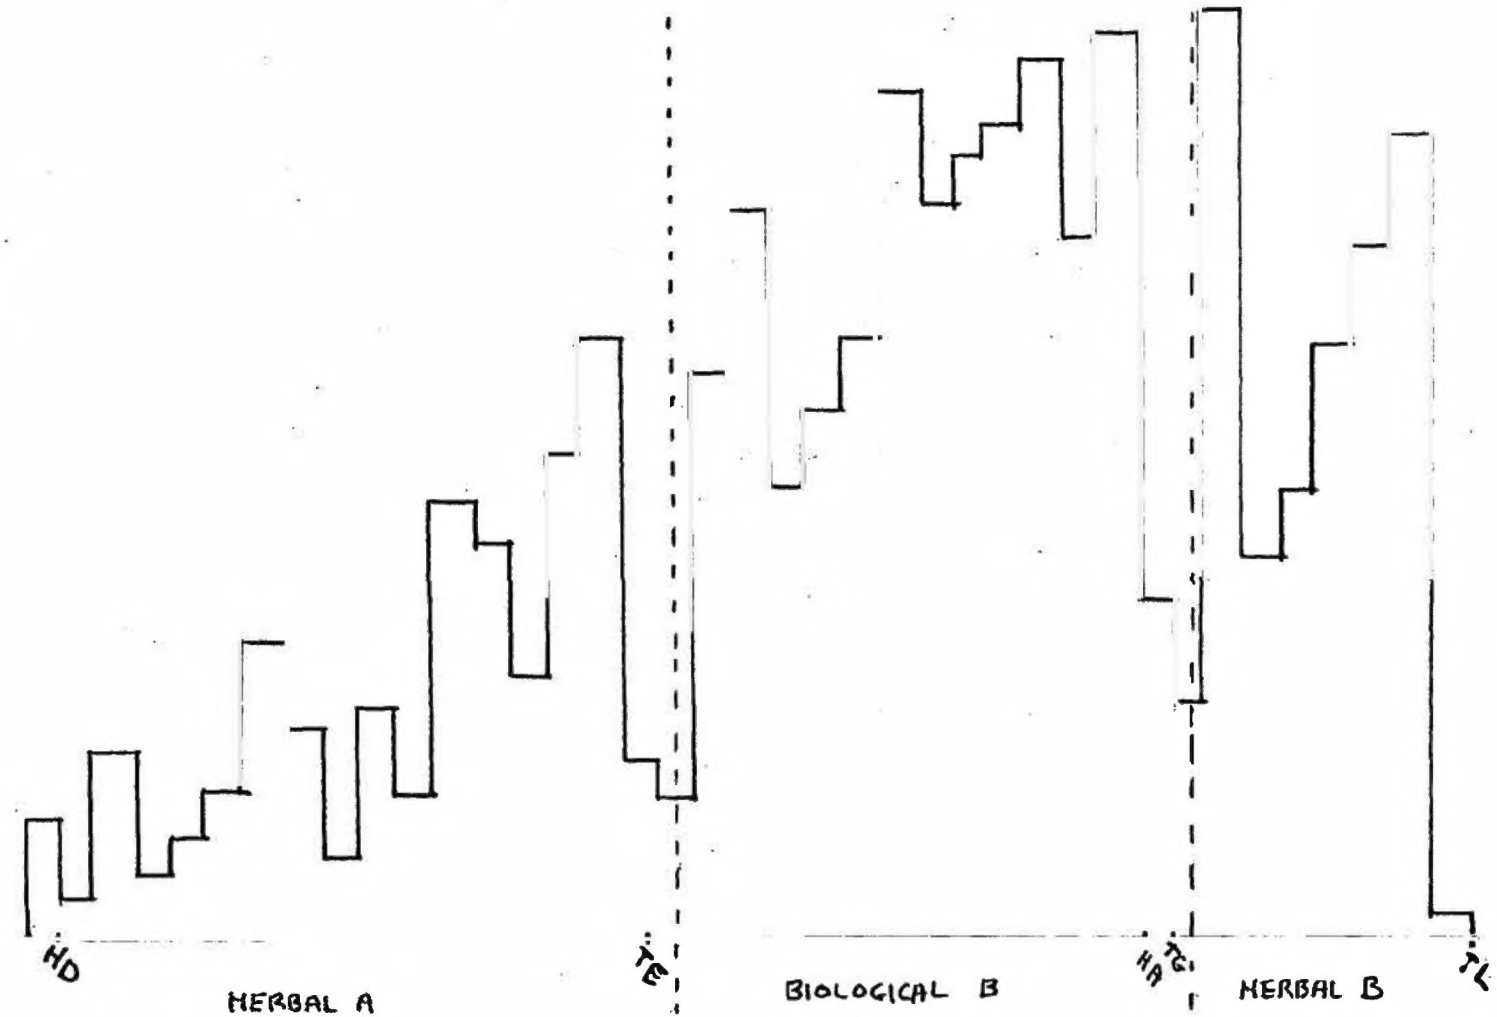

RESULTS OF HICLUS

Figure 4

Doc ID: 6588666

Figure 5

Summary of all Program Results

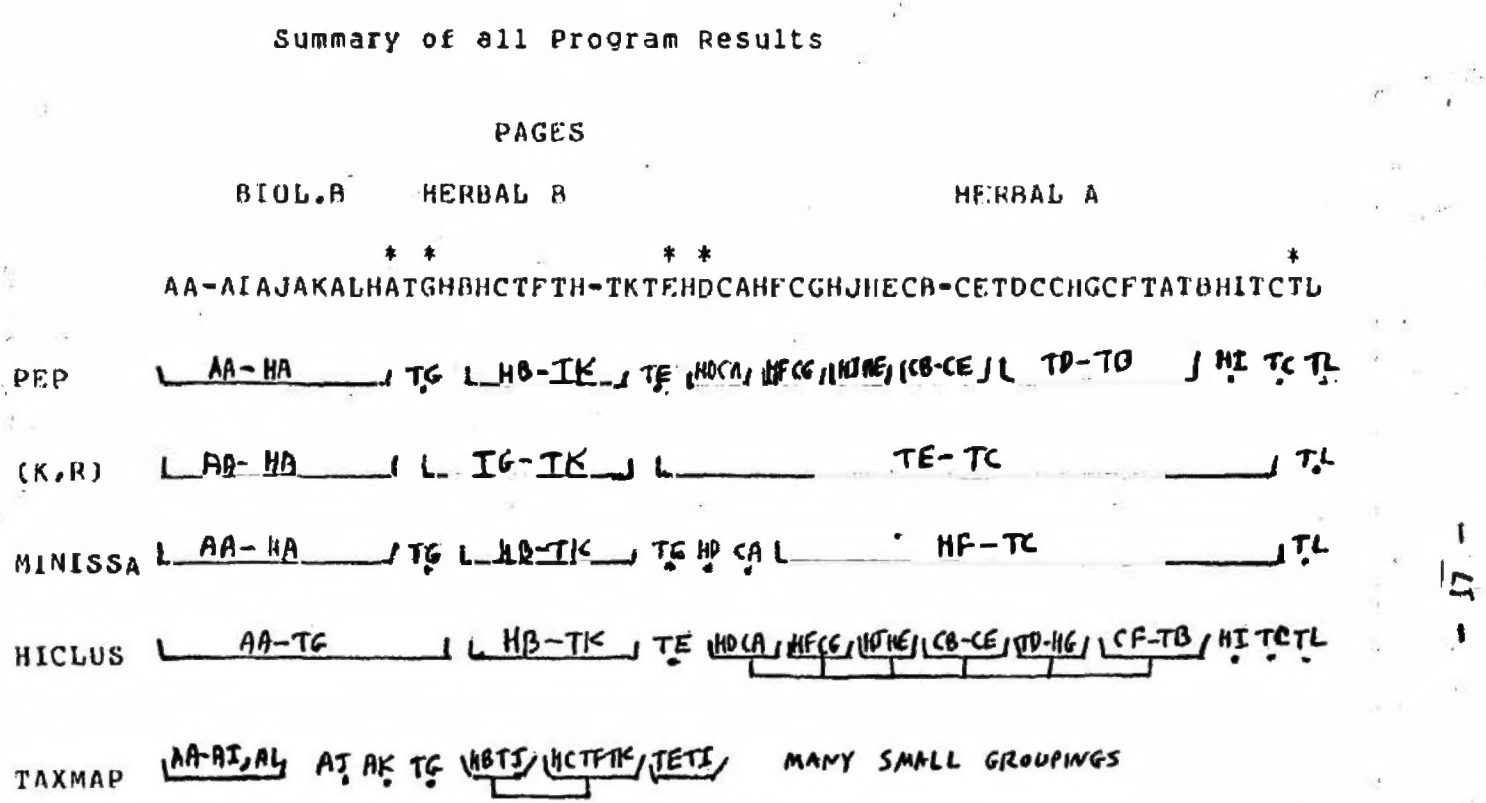

* indicates Herbal B pages which are sometimes separated from other Herbal B

Figure 6

Summary of Samples from Manuscript Pages

no.	code	Currier's page no.	"subject"	"hand"	"language"
1	AA	147			
2	AB	156			
3	AC	149			
4	AD	151			
5	AE	148			
6	AF	150	Biological	Z	B
7	AG	152			
8	AH	153			
9	AI	154			
10	AJ	155			
11	AK	157			
12	AL	158			
13	CA	001			
14	CB	005	Herbal	A	A
15	CC	015			
16	CD	032			
17	CE	045			
18	CF	039			
19	CG	041			
20	HA	059			
21	HB	075			
22	HC	065	Herbal	B	B
23	HD	076			
24	HE	068			
25	HF	072			
26	HG	095			
27	HI	110			
28	HJ	081			
29	TA	082	Herbal	A	A
30	TB	057			
31	TC	109			
32	TD	096			
33	TE	060			
34	TF	066			
35	TG	079			
36	TH	083			
37	TI	084	Herbal	B	B
38	TJ	089			
39	TK	090			
40	TL	094			